UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/732,003	12/07/2000	Rodney A. DeKoning	99-284	1605

24319 7590 02/27/2004

LSI LOGIC CORPORATION
1621 BARBER LANE
MS: D-106 LEGAL
MILPITAS, CA 95035

EXAMINER	
THOMAS, SHANE M	
ART UNIT	**PAPER NUMBER**
2186	6

DATE MAILED: 02/27/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

PTO-90C (Rev. 10/03)

Office Action Summary	Application No.	Applicant(s)
	09/732,003	DEKONING ET AL.
	Examiner	Art Unit
	Shane M Thomas	2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☐ Responsive to communication(s) filed on _____.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) 1-15 is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) 1-15 is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on 08 January 2004 is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: _____.

DETAILED ACTION

Drawings

The Examiner respectfully withdraws the objection under 37 CFR 1.84(p)(5) regarding figure 5.

Claim Rejections - 35 USC § 112

The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

The Examiner respectfully withdraws the rejections of claims 1-7,14, and 15 under §112, second paragraph.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless -
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

The Examiner respectfully withdraws the rejections of claims 1, 2, 5-7, and 13 under §102(e) as anticipated by Nguyen et al. (U.S. Patent Application Publication No. US

2002/0004883) and claims 1, 2, 5-8, and 11-13, also under §102(e), as anticipated by Ito et al. (U.S. Patent No. 6,408,359).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Amended claims **1, 2, 5-8, and 11-13** are rejected under 35 U.S.C. 103(a) as being unpatentable over Otterness et al. (U.S. Patent No. 6,654,831) in view of Nguyen et al. (U.S. Patent Application Publication No. US 2002/0004883).

As per amended **claim 1**, Otterness shows in figure 5 a consolidated storage array that comprises a plurality of storage arrays (combination of controllers 302 with their respective plurality of hard disks 358). As per the amended portion of claim 1 (lines 8-9 and 22-24) Otterness teaches these limitations in column 8, lines 8-40. Specifically, Otterness teaches that any of the storage arrays can be a primary device (lines 22-23) and any of the storage arrays can take over as a primary device for a failed controller of a primary device (lines 38-40). Otterness does not teach a method as described in amended claim 1 (line 22 of claim 1); however, Nguyen does, as will be stated below. Nguyen method would have allowed the consolidated storage array of Otterness to have been able to create virtual data volumes across the data spans, which is only hinted at by Otterness in column 4, lines 23-42.

Therefore it would have been obvious to one having ordinary skill in the art at the time the invention was made to have combined the consolidated storage array system of Otterness, shown in figure 5, with the teachings of Nguyen in order to have been able to change the storage bandwidth on demand by striping data to multiple storage arrays (see the abstract of Nguyen). Thus the teachings of Nguyen would have allowed the CSA of Otterness to have been able to detect when a storage request's needs exceed performance capabilities of a single device (such as a data span 358), and as a result, attempt to allocate more storage devices to stripe the data into multiple storage arrays, thereby meeting the bandwidth requirements of the request (¶39 of Nguyen).

It is important to note that while the disclosure of Nguyen refers to a tape system for describing the invention, it could have been seen by one with ordinary skill in the art at the time the invention was made that hard disk drives could have been used in lieu of tape drives. See ¶41 of Nguyen.

A host device 350 communicates with the storage arrays (combination of a controller 302 and respective data span 358) to perform read and write to a virtual data volume (or LUN, column 5, lines 27-29). The host connects to the consolidated storage array via host loop 352.

As can been seen in the example of paragraph 27, the NSM (network service manager) of Nguyen receives needed storage performance requirements the host. The examiner is also considering the size of the virtual data volume to be included in the --performance requirements-- because the NSM would have had to know how much physical space to allocate. Upon receiving the requirements and analyzing the storage arrays' data rate (performance capacities of data spans 358 of modified Otterness) individually, the NSM

constructs a virtual device (or --logical data volume--) by selecting some of the storage arrays so that the combination of arrays meets the required data rate (performance capacity). Further in paragraph 27, the NSM would have then instructed the storage arrays to stripe the data to storage arrays (302 + 358) as needed, thus achieving the required data rate. The virtualization or creation of the virtual device – herein referred to as a virtual data volume – is performed by functions in the Connection Blocks (CB1-CB9). These functions can be realized in *software* or hardware (see paragraph 28). The Examiner is considering Connection Blocks CB5-CB9 to be part of the data management system of each controller 302, as shown in figure 6 of Otterness and blocks CB1-CB4 contained on each host, respectively.

Referring to ¶29 of Nguyen, if the host device would have requested access to a virtual data volume, the host would have sent a request to the RM software of the NSM, which in turn would have communicated with the ACSLS (maintaining the volume database) to determine if the volume had already been created. If new, the RM software would have made a request to the RA software to reserve the necessary number of storage arrays, and then instructed the RC software to configure the connection blocks (CB5-CB9) for data transfer from the host 350. The RC software also would have instructed the host that the data will be striped to the allocated number of storage arrays. Once the RM software would have been informed that everything is set up, it would have sent a notification to the host that the virtual device (--logical data volume--) is ready for data transfer. Because the RA software would have decided which physical devices to utilize in creating the virtual device and to what locations of the physical device(s) the data would have been sent (and striped if necessary), it is inherent that the information (--volume information--) regarding the physical location of the virtual data volume would have been

contained as an entry in a database apparatus. (Nguyen states in paragraph 35 that when the RM software receives the request from the host, that it ascertains the physical location of the physical medium(s) that contains the selected data volume – that is if the host wishes to read data already contained in the storage array.) Further, it is inherent that if striping data across the storage arrays would have been required, that a --striping definition-- would have accompanied the volume information – either directly or indirectly (pointer) – since the physical location of a virtual data volume could have been ascertained from the consolidated storage array (¶ 35 of Nguyen). Finally (returning to ¶29), the consolidated storage array then sends this volume information to the requesting host device to configure the striping software to write the data in the virtual data volume.

As per **claim 2**, ¶29 (as previously discussed) cites an example of the host device (via data processor DP1) issuing a volume request specifying the data transfer rate, data access time, data compression, etc. to the RM software of the --consolidated storage array--. Volume information regarding which storage arrays and locations are sent to the host. The data access functions (that can be performed in software – paragraph 28) that the Examiner has stated above are contained in the controllers 302 are then configured to stripe the data from the host into the allocated storage arrays.

As per **claim 5**, Nguyen states in paragraph 14 and 37 that the --consolidated storage array-- dynamically configures the storage devices (of the storage array) responding to the requests of the host. These requests can come as manual operator commands, pre-programmed algorithms, rules, application program initiated requests, and the like. Further, in paragraph 15,

Nguyen states the rules that can be implemented comprise response time constraints, data file transfer size, file transfer rates, and data file size bounds.

As per **claim 6**, Nguyen states in paragraph 28 that the [data access] functions in the host can be implemented in software. These functions achieve the virtualization process, or creation of the --logical data volume-- by means of striping the data from the host into the allocated storage arrays. The connection blocks are responsible for the striping of the data to the corresponding storage arrays as cited in paragraph 29.

As per **claim 7**, as has been discussed, data is striped from the processors into a virtual data volume of the storage arrays. Further, the examiner is stating that a --striping definition--, either directly or indirectly, is inherently included with the --volume information-- of the virtual data volume that is stored in the ACSLS or another database. Nguyen's storage system must have a way of accessing the data that has been striped to the storage arrays as well as to determine which location contains the parity information for the stripe.

As per **claim 8**, Otterness shows a plurality of storage array in figure 5, as discussed above. Otterness shows in figure 5 that the storage arrays comprise a plurality of storage devices 356. The rejection for **claim 8** follows **claim 1**'s rejection. The examiner is regarding the --storage area network-- to be the entire apparatus of Otterness as shown in figure 5. The examiner is regarding any of the storage arrays to be a primary device since every storage array is connected to the other storage arrays. Each storage array (modified by the teachings of Nguyen) would have executed volume create software based on requests from host 350 and have analyzed the storage array data rates individually. Further, each [modified] storage array would have constructed (if necessary) a virtual device to meet the data rate performance requirement of

the requesting host (¶27), and would have configured the virtual data volume to accept and distribute data from the host throughout the storage arrays, as taught by Nguyen.

As per **claim 11**, Nguyen states in paragraph 14 that the RC software of the NSM configures the data storage resources [during the creating of a virtual data volume] and responds to manual operator (user) commands, pre-programmed algorithms, rules, application program initiated requests and more. Further, in paragraph 15, Nguyen states that the rules can comprise response time constraints, data file transfer size, data transfer rates, data transfer bounds, and more.

As per **claim 12** with regard to Nguyen, the same rejection as **claim 7** applies. Additionally, it could have been seen that the striping software of Nguyen, contained in the function of the Connection Blocks (refer to paragraph 28), could have been exchanged with the striping and redundancy techniques admitted as prior art by the applicant (page 2, lines 18-19). Since the striping of Nguyen would have taken place within the requesting host's Connection Block (refer to paragraph 29), the requesting processor would not need to realize the physical to virtual translation. Thus, striping data from the host across the storage devices in different storage arrays could have been realized.

As per amended **claim 13**, the examiner is referring to the entire apparatus of figure 5 of Otterness as a --storage area network--. The same rejection for claim 1 is applied to lines 1-6 since the rejection discusses obtaining the performance requirements for the virtual data volume (--logical data volume--), and that if the performance requirement exceeds the capacity of an individual array (in the example cited by Nguyen, data rate was the performance requirement), multiple storage arrays could be combined and the date striped to obtain the required data rate.

The host device --establishes the parameters-- of the virtual data volume in its request to the --consolidated storage array--.

The same rejection for **claim 2** is applied to **lines 7-9 and 12-13 of amended claim 13**.

The same rejection for amended **claim 1** and **claim 7** are applied to lines **13-16 of claim 13**. Claim 1's rejection discusses creating an entry in a database apparatus that contains --volume information-- regarding where the physical location of the data volume is located on the physical medium(s) and that the --striping definition-- is inherently included in the --volume information-- for accessing the data that has been striped to the storage arrays as well as to determine which location contains the parity information for the stripe.

As per **lines 17-18 of amended claim 13**, as was previously discussed in claim 1's rejection, the striping software in the host (contained in functions in the Connection Blocks) uses the --volume information--. Once the RC software sets up the data system shown in figure 5 to transfer data, it instructs the Connection Block, corresponding to the requesting processor, to stripe the data to the allocated virtual data volume contained on storage arrays that the RA software has allocated.

Alternatively, as **per claims 1 (lines 16-20), 7, and 13**, if the applicant can show that the describing of a volume information in a database apparatus is not anticipated by Nguyen, then it would have been obvious to one of ordinary skill in the art at the time the invention was made to create and store volume information, when allocating a virtual data volume, in a database in the NSM, so that after the data had been written (striped if necessary) by the Connection Blocks, it could have been read out correctly when an access to the virtual data volume was made. Further it would have been obvious to include a striping definition – either directly or indirectly (pointer)

– with the volume information in order to allow the processor that had requested the data from the virtual data volume to have had known where the parity information for the data would have been located so that the data could have been verified before being read out.

Claims 3 and 4 are rejected under 35 U.S.C. 103(a) as being unpatentable over Otterness et al. (U.S. Patent No. 6,654,831) in view of Nguyen (U.S. Patent Application Publication No. US 2002/0004883) in further view of Allen et al. (U.S. Patent No. 5,151,990). Nguyen does not utilize a method that monitors (claim 3) and migrates (claim 4) --storage spaces-- of a part of a virtual data volume of a storage array.

As per **claim 3**, Allen describes in column 4, lines 25-40, a method for comparing the remaining capacity of a volume to a predetermined threshold and signaling an error message if the remaining capacity falls below the threshold. The examiner is regarding the size of the storage array to be the --maximum performance capability-- of the array.

As per **claim 4**, Allen describes in column 4, lines 40-48, utilizing the error message to initiate a dynamic reallocation or --data migration-- of available space within the system. This feature allows for operations that would normally have to abort to continue running between the time when the error was generated and the end of the operation.

Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to incorporate the monitoring and reallocation of resources system of Allen onto the Network Service Manager (NSM) as taught by Nguyen, which is part of the consolidated storage array system of modified Otterness, in order to have prevented an overflow situation where the data being striped from the host to the storage arrays would not have been

lost due to a shortage of volume capacity. Referring to claim 4, the examiner is regarding the --first and second-- storage arrays to be any two of storage arrays (figure 5 of Otterness, elements 302+358) that the RM software has allocated for a request by a host (assuming the host requires a performance rate greater than an individual storage array).

Claims 9 and 14 are rejected under 35 U.S.C. 103(a) as being unpatentable over Otterness et al. (U.S. Patent No. 6,654,831) in view of Nguyen (U.S. Patent Application Publication No. US 2002/0004883) in further view of Burgess et al (U.S. Patent No. 5,796,633).

As per **claim 9**, the Examiner is regarding that the --performance capability-- to be either the size (amount of data that can be stored on the array) or the frequency of access (reads/writes) to the storage device, and the --maximum performance capacity-- to be the maximum amount of data that can be stored in the array or the maximum frequency of access the array can withstand. Modified Otterness does not utilize a method that monitors the performance of the storage arrays to determine whether the arrays are performing within a predetermined range of the maximum performance capacity. Burgess shows performance monitoring relating to different performance parameters or thresholds of physical and logical volumes such as frequency of reads and writes (column 8, lines 66-67). These predetermined thresholds can be determined by the user (column, 7, lines 1-3). The user can instruct a configuration file to monitor current memory loads, available physical memory, available page files, and the available virtual memory size (column 6, line 4-9). Further the monitoring system of Burgess can alert the user or a log file if a monitored performance counter dips below its predetermined threshold or --maximum performance capability--. An alert thread can monitor the percentage of free space remaining on each logical volume (column 7, lines 9-16). The monitoring and warning system of Burgess

would determine if a logical volume (storage array) was reaching its maximum storage capacity or the frequency of accesses (read/write) operations. The alert signal could prevent the system from saving further data to the array and be used to trigger a response or action from the system operator to overcome the problem (column 2, lines 45-49). Burguss further states that the functions of the monitoring method could be realized in software (column 15, lines 3-6).

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to further alter the data storage system of Otterness to utilize the monitoring method of Burgess by incorporating the method as a further software piece in the controllers 302 of the storage arrays. This alteration would have allowed for alerts to be flagged when the logical volume capacities of the storage arrays exceeded a threshold determined by the system operator. These alerts would have allowed an operator to give immediate attention to the performance problem. The monitoring of logical volume capacity would have helped to warn of impeding overflow situations where the data being striped from the host to the storage arrays could be lost due to a shortage of volume capacity. Burguss' monitoring method is further beneficial to the system of modified Otterness because the performance monitoring can be handled automatically without human intervention (column 2, lines 35-39).

As per **claim 14**, the same rejection for claim 9 applies. The examiner is regarding the --data transfer performance parameter-- of the storage arrays being monitored as being the frequency of the accesses (reads/writes).

Claims 10 and 15 are rejected under 35 U.S.C. 103(a) as being unpatentable over Otterness et al. (U.S. Patent No. 6,654,831) in view of Nguyen (U.S. Patent Application Publication No. US 2002/0004883) in view of Burgess et al. (U.S. Patent No. 5,796,633) in further view of Komachiya et al. (U.S. Patent No. 6,571,314).

As per **claims 10 and 15**, modified Otterness does not disclose a means or method to allow a system operator or the data system itself (figure 5 of Otterness) to compensate if an alert had been raised signifying that a storage array is within its threshold and maximum capacity or capability. Komachiya teaches optimization of a storage system, where the frequency of accesses to a storage area is over a predetermined threshold, in order to improve the capacity efficiency of the system by migrating the data contained in a single storage area over multiple storage areas. Further, Komachiya teaches that the system can be re-optimized when the frequency of access drops below a predetermined threshold to combine data across multiple storage areas into a single storage area (refer to column 4, lines 60-64 and column 6, lines 43-55). Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to combine the modified system of Otterness with the teaching of Komachiya in order to have been able to migrate a virtual data volume contained in a single storage array, that was being frequently accessed, to multiple storage arrays, thereby increasing the --capacity efficiency-- (or data transfer performance requirement) of the virtual data volume and the overall system. Hence it could have been seen, that the monitoring system of modified Otterness, coupled with the teaching of Komachiya, would have been able to migrate data across multiple storage arrays in order to have overcome the performance degradation of the system when the frequency of accesses (reads/writes) would have exceeded a predetermined threshold,

as would have been defined by the system operator. Here the examiner is referring to the storage array performing near its performance threshold as the --first storage array-- and any storage array that data migrates to as a result to be the --second storage array--. Also, the examiner is regarding the --size-- (number of storage locations) of the virtual data volume to be a --performance capability-- of the volume. The modified system of Otterness would have taken into account the size of the virtual data volume in light of the available size of a --second-- storage array when migrating the volume since the system would obviously not have migrated data to a storage array with an inadequate amount of available storage locations.

As stated in the rejection for claim 14 with regard to Otterness, the examiner is regarding the --data transfer performance parameter-- being monitored as the amount of data loaded (transferred) into the storage array (physical space used) or the frequency of the accesses (reads/writes).

The addition of the teaching of Komachiya would have been obvious because it would have allowed the system to migrate data from a storage array that was close to its maximum performance threshold immediately once the data transfer (access frequency) was detected over the predetermined threshold of the storage array. Thus, coupled with the automatic alert monitoring system of Burgess and data migration teaching of Komachiya, the maintenance burden of the system operator would be significantly reduced.

Amended **claims 1, 2, 5-8, and 11-13** are rejected under 35 U.S.C. 103(a) as being unpatentable over Ito et al. (U.S. Patent No. 6,408,359) in view of Otterness et al. (U.S. Patent No. 6,654,831).

As per **claims 1 and 8**, the examiner is referring to a --consolidated storage array-- as the system of figure 1 combined with the --storage arrays--, the storage arrays comprising storage devices 1807a-1807c with controller 1806a and storage devices 1807d-1807f, respectively, of figure 5. Virtual data volumes (or --logical data volumes--) are distributed equally to each storage array and also to each storage device within the array (column 3, lines 40-53). The examiner is referring to a --host device-- as the external device, which can run software applications for a Video-On-Demand service (column1, lines 21-24), which sends the file creating request to the Destination Determining Part 1803 (figure 1) of the consolidated storage array. The external device sends a --file-- (or virtual data volume) request that includes information specifying the number (size) of storage locations of the virtual data volume to the Destination Determination Part (1803 of figure 1) of the consolidated storage array (column 3, lines 43-49). The examiner will refer to this information as the --performance requirement-- of the virtual data volume. As mentioned above, Ito states that the virtual data volume is spread equally throughout the storage arrays and storage devices so that the consolidated storage array can respond to a read request at a high speed, since access requests do not concentrate on specific storage devices nor specific storage array. Therefore, the performance capability (bandwidth and/or physical volume size) of a single array is less than that of the collective storage arrays comprising the consolidated storage array. Further, it can then been seen that the virtual data

volumes' performance requirements exceed the performance capabilities of a single storage array because they span across the storage arrays in the system of Ito.

Ito lacks the ability to have a controller 1806 (figure 5) resume the responsibilities of writing and reading data of another controller if that controller fails or malfunctions. Otterness shows a system in figure 5 that equips redundant controllers 302 to a plurality of storage devices (data spans 358), and each controller being able to access each other controller's plurality of storage devices upon failure of a controller (column 8, lines 22-29). Therefore it would have been obvious to one having ordinary skill in the art at the time the invention was made to have combined the consolidated storage array system as taught by Ito with the teaching of multiple redundant --master-- controllers as taught by Otterness in order to have been able to have been able to re-direct an I/O request directed to a failed controller to an active controller (column 8, lines 28-29) and further to have been able to access a failed controller's storage devices (data span 358). Such a configuration would have allowed the storage array (combination of controller 1806b with storage devices 1807d-f, for example) of the consolidated storage array of Ito (figure 5) to have accessed the storage devices 1807a-c when a failure or malfunctions is detected with controller 1806a. Such a configuration would have increased data reliability and to have allowed any controller to fail in the consolidated storage array without data loss (column 8, lines 36-38 of Otterness).

The Examiner is considering the combination of a controller 302 with its respective disk span 358 as being the equivalent to the storage array structure of Ito. Figure 6 of Otterness shows the design of controllers 302 of figure 5. As can be seen, Otterness includes the data management software portions of figure 1 of Ito (elements 1801-1803) in Otterness' controller

302 - figure 6 (elements included in RAM memory 332). Further, Otterness includes an I/O interface (1808, figure 5 of Otterness) in each controller 302 (element 310 of figure 6 of Otterness). Thus it could have been seen with the teachings of Otterness that each --storage array-- of modified Ito would have had included the data management elements of figure 1 of Ito and the I/O interface 1808 of Ito's figure 5.

Regarding amended claim 1, lines 8-9, and amended claim 8, line 14, Otterness states in column 8, lines 9-40, that any active controller (of a --storage array--) of figure 5 can be a --primary device-- and can transfer data, thus the modified controllers 1806 of Ito could have performed the same with the applied teaching by Otterness.

Regarding amended claim 1, line 22, as discussed above, the Examiner is considering the combination of the modified controller of Ito 1806 (which with the teaching of Otterness includes I/O interface 1808 and data management software (1801-1803 of figure 1 of Ito)) and the storage devices 1807 to be a --storage array--; thus, the method of amended claim 1, as taught by Ito, could have been performed by the controller 1806 of an active (primary) device of modified Ito.

Regarding amended claim 1, lines 23-34, and amended claim 8, lines 27-28, Otterness states in column 8, lines 39-41, that any --storage array-- can become an active [primary] device if a [primary] device fails and take over for the failed device.

Modified Ito's system indirectly would have analyzed the storage arrays for their performance capabilities (physical volume size, or remaining volume size in this instance) by keeping track of empty storage locations in the Empty Area Managing Part (1802 in figure 1). This part manages the empty --blocks-- (storage locations not containing data associated with a

virtual data volume) of each of the storage devices using an address of each empty storage location as well as first and second identifiers, which identify a specific storage device and specific storage array, respectively (refer to column 3, lines 19-39).

The Destination Determination Part (1803 of figure 1) selects and configures locations from each of the storage devices of each storage array by utilizing the information provided by the Empty Area Managing Part and the Configuration Managing Part (1801 of figure 1), which manages the number of storage devices and respective identifying information of the system. The flow chart of figure 6 shows a method used to select and configure empty storage locations by selecting a set number of locations per loop cycle (steps S2305-S2307), allocating them for the virtual data volume (step S2308), comparing the number of locations allocated to the number requested in the performance requirement for the volume sent by the host (step S2309), and selecting and configuring more locations if the number allocated is less than the number requested. Thus the method of figure 6 shows how the performance requirements (size of volume requested by the host) are met by the consolidated storage area system in figure 5 (see column 3, lines 19-39).

Volume information would have been created and managed by the system in the address position file such like element 2401 of figure 7. Each virtual data volume has an address position file associated with it (column 11, lines 11-17). It is inherent that once the virtual data volume is created among the storage arrays and the storage devices that the applications running on a host device, which sent the volume request to the system of Ito, would have used the allocated storage locations to store and retrieve data. Address position file (volume information) is sent from the consolidated storage array to the host device. The I/O control unit 1808

distributes the data inputted from the external device of the host to the storage device control units (1806a and 1806b). These control units then write the data to the allocated storage locations (refer to column 8, lines 54-61). As discussed above, the Examiner is considering the I/O control unit 1808 (interface) and the storage device control units to be comprised in the controllers 302 of the storage arrays as shown in figure 5 of Otterness.

Regarding **claim 8**, the Configuration Managing Part 1801 of figure 1 stores the number of storage devices and respective identifying information, and as has been mentioned above, is being considered by the Examiner to be integrated into the controllers of the storage arrays, as taught by Otterness in figure 6. Thus it could have been seen that a --CSA primary device-- (one of the storage arrays, is connected to *all other* the storage arrays as shown in figure 5 of Otterness. Finally, the examiner refers to the collection of the host, consolidated storage array, and the CSA primary device as a --storage area network--.

As per **claim 2**, as has been shown in the rejection for claims 1 and 8 above of Ito, an external device (of a host) issues a file (or virtual volume) create command to the Destination Determining Part of the --consolidated storage array--. Once the volume has been allocated, the CSA sends the address position file (volume information) to the host device to enable the applications running on the host to utilize the allocated volume to store and retrieve data.

As per **claims 5 and 11**, as has been shown in the rejection for claims 1 and 8 above of modified Ito, the external system specifies the size requirement for the virtual data volume by the number of storage blocks (locations) it needs in order to sufficiently store its data (see column 5, lines 43-49). Further, as has been stated, the examiner is referring to the --size-- of the volume to be the --performance requirement-- of the volume.

Regarding **claim 11**, Ito is referring to a user as something that is utilizing the storage device management system to access (store/retrieve) data. In one example, the system of modified Ito is being used in a Video-On-Demand system as a video server. In this case, it could have been seen that the --user-- could be a --client-- of the server and is requesting a video object and would therefore read data from modified Ito's system. More indirectly, the Examiner is interpreting a user to be the person controlling the --client-- wishing to access the --videos-- contained on the system of modified Ito.

As per **claim 6**, as has been detailed in the rejection for claims 1 and 8 above, the Destination Determining Part of the --consolidated storage array-- allocates storage locations within each storage device of each storage array for each request for a virtual data volume. The software application running on the external device of the host sends data through the external device and supplies it to the I/O control unit (1808 of figure 5) which then sends the data, using the address position file (volume information) received from the consolidated storage array, to the respective storage device control units (1806a and 1806b). The examiner is referring to this procedure as --striping-- data from the host device since the data is not written a specific storage device nor specific storage array.

As per claim **7 and 12**, the same definition of striping as defined in the rejection for claim 6 is herein used for claim 7 and 12's rejection. Further, a --striping definition-- in incorporated into the address position file (volume information) that is created and sent from the Destination Determining Part (1803 of figure 1) of the consolidated storage array. This --striping definition-- supplies the host's I/O control unit 1808 with the address of the allocated storage locations of the

virtual storage array so it can begin sending the data received from the host's external device to the respective storage control units 1806, which write the data to the storage devices 1807.

As per amended **claim 13**, the same rejection for amended claim 1 is applied to lines 1-6; the same rejection for claim 2 is applied to lines 7-9 and 12-13; and the same rejection of claims 6 and 7 are applied to lines 13-18. Further, the rejection of amended claim 1, lines 8-9, is applied to lines 10-11, and the rejection of amended claim 1, lines 22-24, is applied to lines 19-22.

Claims 3, 4, 9, 10, 14, and 15 are rejected under 35 U.S.C. 103(a) as being unpatentable over Ito (U.S. Patent No. 6,408,359) in view of Otterness et al. (U.S. Patent No. 6,654,831) in further view of D'Errico et al. (U.S. Patent No. 6,314,503).

As per **claims 3, 4, 9, 14, and 15**, Ito states that although data is distributed to all storage devices in each storage array of the embodiment described in the above rejections, an alternate embodiment could have data distributed to specific storage devices in each storage array (column 11, lines 65-67). Ito does not disclose in either embodiment a method for monitoring the data transfer performance, which the examiner is referring to as a --performance capability--, of the storage arrays in order to determine whether the storage arrays are performing within a predetermined range of the maximum --performance capability-- of each array. Additionally, Ito does not disclose a method to alleviate the condition of an array performing within its predetermined maximum capability (or maximum data transfer) by migrating a portion of the data between the over-performing (--first--) storage array and another (--second--) storage array.

D'Errico teaches in column 3, lines 49-63, that in a system with multiple storage devices

a performance condition can be detected (hence monitored) and alleviated by the re-distribution of system data between the multiple storage devices. Specifically, D'Errico teaches that in a system with a plurality of storage devices, a method comprises the steps of: (A) detecting a segment in the storage system that is accessed frequently and sequentially (a virtual volume) and is stored on one of the plurality of storage devices and (B) in response to step (A), splitting the large data segment into at least two smaller data segments that can be accessed in parallel from at least two of the plurality of storage devices, thereby improving the performance of the overall system.

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate the teaching of D'Errico to the data storage system of Ito, in order to alleviate the burden of an array that is performing within its predetermined maximum performance capability by reallocating the virtual data volume that is causing the increase in the data transfer performance of the array, by migrating a portion of the virtual data volume to at least a second storage array. Specifically D'Errico's method should have been incorporated into the consolidated storage array's --storage arrays-- (combination of modified Ito's controllers 1806 with their respective storage devices 1807) since it performs all of the managing procedures of the storage arrays and virtual data volumes, as has been described above. The teaching of D'Errico would have allowed the modified system of Ito to migrate frequently accessed data (data with a detected high transfer performance) and have distributed it among other storage arrays and devices; thus, an increase in the amount of data that could have been accessed by an application using the data storage system of modified Ito in the same amount of time would have been achieved. Further, the splitting of the virtual data volume among the storage arrays and

devices would be transparent to the host (and therefore the application running on the host). This aspect is advantageous since the method of D'Errico could have been performed automatically without requiring modification to the application running on the host device, and without requiring manual intervention [by the a system operator] (refer to column 6, lines 60-67 thru column 7, line 15). Therefore, it could have been seen that the teaching of D'Errico would have reduced the burden of the data system by increasing the rate at which data could have been accessed, thus improving system performance by splitting a virtual data volume among storage arrays and storage drives while the division would have remained transparent to the host that would have utilized the data system of Ito.

Regarding amended **claim 14**, the rejections of lines 13-16, follow the rejection of amended claim 1, lines 22-24.

As per **claim 10**, the Examiner is referring to the maximum data transfer performance as the --maximum capacity-- of a storage array. Further, the Examiner is considering the performance capability of the storage arrays to be the size of the storage arrays and the performance requirement of the --logical data volume-- (virtual data volume) to be the size (number of storage locations) of the volume. Therefore, it would have been obvious to one of ordinary skill in the art that when the predetermined range of the data transfer performance had been crossed for a --first-- storage array, the teaching of D'Errico would have compared the sizes of the remaining storage arrays in light of the size of the virtual data volume in order to find a --second-- (or more) storage array(s) with enough empty (unallocated) storage locations to accommodate the portion of the virtual data partition that had been split.

Response to Amendment

As per amended claims 1,8,13, and 14 the Applicant asserts that Nguyen nor Ito teach

(a) that the storage arrays *themselves* perform the necessary functions as claimed;

(b) that the any of the storage arrays can perform the functions as claimed; and

(c) that any of the storage arrays can take over the responsibility of a primary array if failure of the primary array is detected.

The Examiner agrees; however, in response to amended claims **1,8,13, and 14,** the combination of *Nguyen* and *Otterness* and the combination of *Ito and Otterness* teach all of the claimed subject matter of the aforementioned claims.

Response to Arguments

Applicant's arguments, see page 8-10 of Paper No. 5, filed 4 January 2004, with respect to the §112, second paragraph rejection of claims 4, 14, and 15 have been fully considered and are persuasive. The §112, second paragraph rejections of claims 4,14, and 15 have been withdrawn.

Applicant's arguments with respect to claims 1-15 have been considered but are moot in view of the new ground(s) of rejection.

As per amended claims 1,8,13, and 14, the Examiner has applied *Otterness* to reject the claims.

Conclusion

Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Shane M Thomas whose telephone number is (703) 605-0725. The examiner can normally be reached on M-F 8:30 - 5:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matt M Kim can be reached on (703) 305-3821. The fax phone numbers for the organization where this application or proceeding is assigned are (703) 764-7239 for regular communications and (703) 764-7239 for After Final communications.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.

SMT

Shane M Thomas
February 25, 2004

MTK

MATTHEW KIM
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 21(